On Bayesian based adaptive confidence sets for linear functionals

Botond Szabó

arXiv:1412.0459v2 [math.ST] 18 Apr 2015

Abstract We consider the problem of constructing Bayesian based confidence sets for linear functionals in the inverse Gaussian white noise model. We work with a scale of Gaussian priors indexed by a regularity hyper-parameter and apply the data-driven (slightly modified) marginal likelihood empirical Bayes method for the choice of this hyper-parameter. We show by theory and simulations that the credible sets constructed by this method have sub-optimal behaviour in general. However, by assuming "self-similarity" the credible sets have rate-adaptive size and optimal coverage. As an application of these results we construct L_∞-credible bands for the true functional parameter with adaptive size and optimal coverage under self-similarity constraint.

1 Introduction

Uncertainty quantification is highly important in statistical inference. Point estimators without confidence statements contain only a limited amount of information. Bayesian techniques provide a natural and computationally advantageous way to quantify uncertainty by producing credible sets, i.e. sets with prescribed (typically 95%) posterior probability. In this paper we investigate the validity of such sets from a frequentist perspective. We are interested whether these sets can indeed be used as confidence sets or by doing so one gives a misleading uncertainty quantification, see for instance [12]. We focus in our work on credible sets for linear functionals in nonparametric models and their application to the construction of L_∞-credible bands for the functional parameter.

In infinite dimensional models nonparametric priors usually have a tuning- or hyper-parameter controlling the fine details of the prior distribution. In most of the cases the choice of the hyper-parameter in the prior distribution is very influential, incorrect choices can result in sub-optimal behaviour of the posterior. Therefore data driven, adaptive Bayesian techniques are applied in practice to determine the value of the hyper-parameter, overcoming overly strong prior assumptions. The two (perhaps) most well-known Bayesian methods used to achieve adaptive results are the hierarchical Bayes and the empirical Bayes techniques. In our work we focus mainly on the empirical Bayes method, but we conjecture results about the hierarchical Bayes approach as well.

The frequentist properties of adaptive Bayesian credible sets were considered only in a limited number of recent papers; see [22, 23, 1, 18, 20]. The authors of these papers have shown that under a relatively mild and natural assumption on the functional parameter, i.e. the self-similarity condition, the credible sets have good frequentist coverage and in a minimax sense optimal size. However, for non self-similar functions the credible sets provide overconfident, misleading confidence statements. In these papers mainly the L_2-norm were considered, which is the natural extension of the finite dimensional Euclidean-norm, but for

Botond Szabó
Budapest university of Technology, Műegyetem rkp. 3-9, H-1111 Budapest, Hungary, e-mail: bszabo@math.bme.hu

visualization purposes it is perhaps not the most appropriate choice. In practice usually the posterior credible bands are plotted, which correspond to the L_∞-norm. The frequentist properties of L_∞-credible bands were investigated in a non-adaptive setting in [6, 24]. Adaptive L_∞-credible bands were only considered up to now in the recent work [18], using techniques developed in [6] and [14]. In our work we also focus on the construction of L_∞-credible bands taking a substantially different and independently developed approach than in [18].

In our analysis we consider a sub-class of (possibly non-continuous) linear functionals satisfying the self-similarity condition and construct credible sets over it using empirical Bayes method on a scale of Gaussian priors with varying regularity; see also [22] for the application of this family of priors to derive L_2 credible sets for the functional parameter. We show that by slightly modifying the empirical Bayes procedure we can construct credible sets with rate adaptive size and good coverage property for the linear functionals of the self-similar functional parameters. However, there exist certain oddly behaving functional parameters (not satisfying the self-similarity assumptions), where the empirical Bayes method provides haphazard and misleading credible sets for the linear functionals of the functional parameter. This result is in itself of independent interest, since until now the frequentist properties of credible sets in semi-parametric problems were mostly investigated in non-adaptive settings, see for instance [2, 5, 19, 15] and references therein.

However, perhaps the main contribution of the present paper is the application of the derived results about linear functionals to the analysis of L_∞-credible bands. We show that point evaluations of the functional parameter satisfy the self-similarity assumption on linear functionals. Therefore the above described results about credible sets for linear functionals apply also to pointwise credible sets for the functional parameter. Then putting together these pointwise credible sets we arrive at an L_∞-credible band, which therefore has again good frequentist properties for self-similar functional parameters. This technique is essentially different from the one applied in [18] for the construction of L_∞-credible bands, where wavelet basis with spike and slab priors were considered and a weak Bernstein-von Mises theorem was proved.

The remainder of the paper is organized as follows. In Section 2 we introduce the inverse Gaussian white noise model, where we have carried out our analysis. Then in Section 2.1 we introduce the linear functionals we are interested in (satisfying the self-similarity constraint) and show that the point evaluations of the functional parameter satisfy this property. The construction of the empirical Bayes credible sets are given in Section 2.2. The main results of the paper are formulated in Sections 2.3 and 2.4. In Section 3 we provide a short numerical analysis demonstrating both the positive and negative findings of the paper. The proofs of the main theorems are deferred to Sections 4 and 5.

2 Main result

Consider the inverse Gaussian white noise model

$$X_t = \int_0^t \mathbb{K}\theta_0(s)ds + \frac{1}{\sqrt{n}}B_t, \qquad t \in [0,1],$$

where B_t is the Brownian motion, $1/n$ is the noise level, X_t the observed signal, $\theta_0(\cdot) \in L^2[0,1]$ the unknown function of interest and $\mathbb{K} : L^2[0,1] \mapsto L^2[0,1]$ a given compact, linear, self-adjoint transformation (but we also allow $\mathbb{K} = \mathbb{I}$). From the self-adjoint property of $\mathbb{K}$ follows that its eigenvectors $\varphi_i(\cdot) : [0,1] \mapsto \mathbb{R}$ form an orthogonal basis and the compactness ensures that the corresponding eigenvalues κ_i are tending to zero. Hence using series expansion with respect to φ_i we get the equivalent Gaussian sequence model

$$X_i = \kappa_i\theta_{0,i} + \frac{1}{\sqrt{n}}Z_i, \quad \text{for all } i = 1,2,... \tag{1}$$

where $X_i = \langle X., \varphi_i(\cdot)\rangle$ and $\theta_{0,i} = \langle \theta_0(\cdot), \varphi_i(\cdot)\rangle$ are the series decomposition coefficients of the observation and the true function, respectively and the random variables $Z_i = \langle B., \varphi_i(\cdot)\rangle$ are independent and standard normal distributed. We limit ourselves to the mildly ill-posed inverse problems, where

$$C^{-2}i^{-2p} \leq \kappa_i^2 \leq C^2 i^{-2p}, \tag{2}$$

with some fixed non-negative constant p and positive C, see [7] for the terminology.

Suppose furthermore that the unknown infinite dimensional parameter $\theta_0 = (\theta_{0,1}, \theta_{0,2}, ..)$ belongs to a hyper-rectangle

$$\Theta^\beta(M) = \{\theta : \theta_i^2 i^{1+2\beta} \leq M \quad \text{for all } i = 1,2,...\} \tag{3}$$

where β is the regularity parameter and M is the squared radius of the hyperrectangle. The minimax estimation rate of the full parameter θ_0 is a multiple of $n^{-\beta/(1+2\beta+2p)}$, see [10].

2.1 Linear functionals

In this paper we focus on the construction of confidence sets for the (possibly unbounded) linear functionals

$$L\theta = \sum l_i \theta_i, \tag{4}$$

where $l = (l_1, l_2, ...)$ is in a self-similar hyper-rectangle $L_s^q(R)$, for some q, R, j_0, K

$$L_s^q(R) = \{l \in \ell_2 : (1/R^2) j^{-1-2q} \leq \sum_{i=j}^{j+K-1} l_i^2 \leq R^2 j^{-1-2q}, \text{for all } j > j_0\}, \tag{5}$$

where the parameters j_0 and K are omitted from the notation. We note that for instance the linear functionals in the form $l_i \asymp i^{-q}$ belong to this hyper-rectangle.

We are particularly interested in the class of non-continuous linear functionals, the point evaluations of the functional parameter θ. For a specific choice of the operator $\mathbb{K}$ all point evaluations on $t \in [0,1]$ belong to $L_s^{-1/2}(R)$; see the next paragraph. Therefore confidence sets for self-similar linear functionals $L \in L_s^{-1/2}(R)$ of the series decomposition coefficients θ_0 also provide us pointwise confidence sets of the function $\theta_0(\cdot)$. Gluing together the (uniform) pointwise confidence sets one arrives to L_∞ confidence bands.

In this paragraph we show that point evaluations of the function $\theta_0(\cdot)$ belong to the self-similar class of linear functionals for appropriate choice of the basis. Assume that the eigen-basis of the operator $\mathbb{K}$ is the sine-cosine basis $\varphi_i(\cdot)$. The function $\theta_0(\cdot)$ can be given with the help of the trigonometric decomposition

$$\theta_0(t) = \sum_i \theta_{0,i} \varphi_i(t). \tag{6}$$

Since $\varphi_{2i+1}^2(t) + \varphi_{2i}^2(t) = \sin^2(i2\pi t) + \cos^2(i2\pi t) = 1$ we got that $l_i = \varphi_i(t)$ is in $L_s^{-1/2}(2)$ with parameters $j_0 = 1$ and $K = 3$ (since every three consecutive integers contain a pair of $(2i-1, 2i)$ for some $i \in \mathbb{N}$).

2.2 Bayesian approach

In the Bayesian framework to make inference about the unknown sequence θ_0 we endow it with a prior distribution. In our analysis we work with the infinite dimensional Gaussian distribution

$$\Pi_\alpha = \bigotimes_{i=1}^{\infty} N(0, i^{-1-2\alpha}), \tag{7}$$

where the parameter $\alpha > 0$ denotes the regularity level of the prior distribution. One can easily compute the corresponding posterior distribution

$$\Pi_\alpha(\cdot\,|\,X) = \bigotimes_{i=1}^{\infty} N\Big(\frac{n\kappa_i^{-1}}{i^{1+2\alpha}\kappa_i^{-2}+n}X_i, \frac{\kappa_i^{-2}}{i^{1+2\alpha}\kappa_i^{-2}+n}\Big). \tag{8}$$

Furthermore by combining and slightly extending the results of [15] and [4] one can see that the choice $\alpha = \beta$ leads to posterior contraction rate $n^{-\frac{\beta}{1+2\beta+2p}}$ for $\theta_0 \in \Theta^\beta(M)$, while other choices of the parameter α provide sub-optimal contraction rates.

In this paper, however, we are interested in the posterior distribution of the linear functional $L\theta$. From Proposition 3.2 of [15] follows that the posterior distribution of the linear functionals $L\theta$ (assuming measurability with respect to the prior Π_α) takes the form

$$\Pi_\alpha^L(\cdot|X) = N\Big(\sum_i \frac{nl_i\kappa_i^{-1}}{i^{-1-2\alpha}\kappa_i^{-2}+n}X_i, \sum_i \frac{l_i^2\kappa_i^{-2}}{i^{1+2\alpha}\kappa_i^{-2}+n}\Big). \tag{9}$$

Furthermore it was also shown in Section 5 of [15] that the optimal choice of the hyper-parameter α is not β, but rather $\beta - 1/2$. The resulting optimal rate is of the order $n^{-(\beta+q)/(2\beta+2p)} \vee n^{-1/2}$; see [9, 11]. Note that in case $q \geq p$ the smoothness of the linear functional compensates for the degree of ill-posedness and we get a regular problem with contraction rate $n^{-1/2}$. However, in our work we focus on the (from the construction of credible bands point of view) more interesting case $q < p$.

Since the regularity parameter β of the infinite sequence θ_0 is usually not available one has to use data-driven method to choose α, which we will refer from now on as the hyper-parameter of the prior. Following [14] and [22] we select a value for α with the marginal likelihood empirical Bayes method, i.e. we select the maximizer of

$$\hat\alpha_n = \arg\max_{\alpha\in[0,\mathrm{A}]} \ell_n(\alpha), \tag{10}$$

where A is some arbitrary large, fixed constant, and ℓ_n denotes the corresponding log-likelihood for α (relative to an infinite product of $N(0,1/n)$-distributions)

$$\ell_n(\alpha) = -\frac{1}{2}\sum_{i=1}^{\infty}\Big(\log\Big(1+\frac{n}{i^{1+2\alpha}\kappa_i^{-2}}\Big) - \frac{n^2}{i^{1+2\alpha}\kappa_i^{-2}+n}X_i^2\Big). \tag{11}$$

Then the *empirical Bayes posterior* for the functional parameter θ is defined as $\Pi_{\hat\alpha_n}(\cdot|X)$ obtained by substituting $\hat\alpha_n$ for α in the posterior distribution (8), i.e.

$$\Pi_{\hat\alpha_n}(B|X) = \Pi_\alpha(B|X)\Big|_{\alpha=\hat\alpha_n}$$

for measurable subsets $B \subset \ell^2$. Slightly adapting the proof of Theorem 2.3 in [14] we can get that the posterior distribution of the functional parameter θ achieves the corresponding minimax contraction rate up to a logarithmic factor.

As conjectured in page 2367 of [15] and Section 2.3 of [14] this suggests that the present procedure is sub-optimal for the linear functional $L\theta_0$, since adaptation for the full parameter θ_0 and its linear functionals $L\theta_0$ is not possible simultaneously in this setting. In view of the findings in the non-adaptive case [15] we might expect, however, that we can slightly alter the procedures to deal with linear functionals. For instance, it is natural to expect that the empirical Bayes posterior for linear functionals $L\theta$, given in (4),

$$\Pi^L_{\hat\alpha_n-1/2}(\cdot\,|\,X) = \Pi^L_\alpha(\cdot\,|\,X)\Big|_{\alpha=\hat\alpha_n-1/2} \tag{12}$$

yields optimal rates. In the present paper we work with this data-driven choice of the hyper-parameter and investigate the frequentist properties of Bayesian credible sets constructed from the posterior (12).

For fixed hyper-parameter α the posterior (9) is a one dimensional Gaussian distribution hence a natural choice of the credible set is the interval

$$\hat{C}_{n,\alpha} = [\widehat{L\theta}_{n,\alpha} - \zeta_{\gamma/2} s_n(\alpha), \widehat{L\theta}_{n,\alpha} + \zeta_{\gamma/2} s_n(\alpha)], \tag{13}$$

where $\widehat{L\theta}_{n,\alpha}$ is the posterior mean, $s_n^2(\alpha)$ the posterior variance given in (9), and ζ_γ is the $1-\gamma$-quantile of the standard normal distribution. We note that the mean $\widehat{L\theta}_{n,\alpha}$ of the posterior distribution of the linear functional is exactly the linear functional L of the posterior mean of the full parameter $\hat{\theta}_{n,\alpha}$ given in (8). One can easily see that the preceding interval accumulates $1-\gamma$ fraction of the posterior mass. Then the empirical Bayes credible sets are obtained by replacing α with the data-driven choice $\hat{\alpha}_n - 1/2$ in (13). We introduce some additional flexibility, by allowing the blow up of the preceding interval with a constant factor $D > 0$, i.e.

$$\hat{C}_n^L(D) = [\widehat{L\theta}_{n,\hat{\alpha}_n-1/2} - D\zeta_{\gamma/2} s_n(\hat{\alpha}_n - 1/2), \widehat{L\theta}_{n,\hat{\alpha}_n-1/2} + D\zeta_{\gamma/2} s_n(\hat{\alpha}_n - 1/2)]. \tag{14}$$

2.3 Negative results

However, matters seem to be more delicate than one would expect from the direct case. For certain oddly behaving functions the posterior distribution (12) achieves only sub-optimal contraction rates. Furthermore the credible sets (14) have also coverage tending to zero.

Theorem 1. *Let n_j be positive integers such that $n_1 \geq 2$ and $n_j \geq n_{j-1}^4$ for every j. and let $K > 0$. Let $\theta_0 = (\theta_{0,1}, \theta_{0,2}, \ldots)$ be such that*

$$\theta_{0,i} = \begin{cases} K n_j^{-\frac{1/2+\beta}{1+2\beta+2p}}, & \text{if } n_j^{\frac{1}{1+2\beta+2p}} \leq i < 2n_j^{\frac{1}{1+2\beta+2p}}, \qquad j = 1, 2, \ldots, \\ 0, & \text{else}, \end{cases} \tag{15}$$

for some positive constants β, p. Then the constant $K > 0$ can be chosen such that the coverage of the credible set tends to zero for every $q \in \mathbb{R}$, $D, R > 0$ and $L \in L_s^q(R)$

$$\mathrm{P}_{\theta_0}(\theta_0 \in \hat{C}_{n_j}(D)) \to 0.$$

Furthermore the posterior distribution attains sub-optimal contraction rate

$$\Pi^L_{\hat{\alpha}_{n_j}-1/2}(L\theta : |L\theta_0 - L\theta| \geq m n_j^{-(\beta+q)/(1+2\beta+2p)} | X) \overset{\mathrm{P}_{\theta_0}}{\to} 1, \tag{16}$$

as $j \to \infty$ for a positive, small enough constant m and linear functional L satisfying (5)*.*

The proof of the Theorem is given in Section 4. The sub-optimal contraction rate of the posterior distribution and the bad coverage property of the credible sets are due to the mismatch of the underlying loss functions. In the empirical Bayes method the hyper-parameter α is chosen to maximize the marginal likelihood function. This method is related to minimizing the Kullback-Leibler divergence between the marginal Bayesian likelihood function and the true likelihood function. At the same time the evaluation of the posterior distribution is given with respect to some linear functional L of the functional parameter θ_0. Optimal contraction rate and good coverage follows from optimal bias-variance trade-off. However, the likelihood based empirical Bayes method intends to minimize the Kullback-Leibler divergence, which is not an appropriate approach in general for balancing out the bias and variance terms. Therefore the empirical Bayes (and we believe that also the hierarchical Bayes method) leads to sub-optimal rate and poor coverage.

2.4 Self-similarity

To solve this problem we can introduce some additional constraint on the regularity class $\Theta^\beta(M)$. The notion of self-similarity originates from the frequentist literature [17, 13, 3, 8, 16], and was adapted in the Bayesian

literature [22, 18, 1]. We call a series in the hyper-rectangle $\theta_0 \in \Theta^\beta(M)$ *self-similar* if it satisfies

$$\sum_{i=N}^{\rho N} \theta_{0,i}^2 \geq \varepsilon M N^{-2\beta}, \quad \forall N \geq N_0, \tag{17}$$

where N_0, ε and ρ are some fixed positive constants. Furthermore we denote the class of functions satisfying the self-similar constraint by $\Theta_s^\beta(M)$, where we omit the parameters N_0, ε, ρ from the notation. We denote by $\Theta_s(M)$ the collection of self-similar functions with regularity in a compact interval of regularity parameters $\beta \in [\beta_{\min}, \beta_{\max}]$

$$\Theta_s(M) = \cup_{\beta \in [\beta_{\min}, \beta_{\max}]} \Theta_s^\beta(M). \tag{18}$$

Here we omit again in the notation the dependence on $\beta_{\min}$ and $\beta_{\max}$ and assume that $\beta_{\min} > -q$, else $L\theta_0$ would be infinite.

We show that uniformly over $L \in L_s^q(R)$ and $\theta_0 \in \Theta_s(M)$ the coverage of credible sets $\hat{C}_n^L(D)$ for the linear functionals $L\theta_0$ tends to one. Furthermore we prove that the size of the credible sets achieves the corresponding minimax contraction rate.

Theorem 2. *There exists a large enough positive constant D such that the empirical Bayes credible sets $\hat{C}_n^L(D)$ have honest asymptotic coverage one over the self-similar linear functionals $L \in L_s^q(R)$ of the functional parameter θ_0 satisfying* (5)*, i.e.*

$$\inf_{\theta_0 \in \Theta_s(M)} P_{\theta_0}\big(L\theta_0 \in \hat{C}_n^L(D), \forall L \in L_s^q(R)\big) \to 1. \tag{19}$$

Furthermore the radius of the credible sets is rate adaptive, i.e. there exists a positive constant $C_1 > 0$ such that for all $\beta \in (q, \beta_{\max}]$ we have

$$\inf_{\theta_0 \in \Theta_s^\beta(M)} \mathrm{P}_{\theta_0}\big(s_n(\hat{\alpha}_n - 1/2) \leq C_1 n^{-\frac{\beta+q}{2\beta+2p}}, \forall L \in L_s^q(R)\big) \to 1. \tag{20}$$

We defer the proof to Section 5. The credible band on $[0,1]$ can be constructed with the help of the linear functionals $\varphi_i(t)$ introduced in (6), i.e. the point evaluations of the basis $\varphi_i(\cdot)$ at $t \in [0,1]$. Following from its definition (14) the credible band takes the form

$$[\hat{\theta}_n(t) - D\zeta_{\gamma/2} s_n(t, \hat{\alpha}_n - 1/2), \hat{\theta}_n(t) + D\zeta_{\gamma/2} s_n(t, \hat{\alpha}_n - 1/2)], \quad t \in [0,1], \tag{21}$$

where $\hat{\theta}_n(t)$ is the posterior mean and $s_n^2(t, \hat{\alpha}_n - 1/2)$ is the posterior variance for $\alpha = \hat{\alpha}_n - 1/2$ given in (12) belonging to the linear functional $L = (l_i)_{i \geq 1} = (\varphi_i(t))_{i \geq 1}$. By combining Theorem 2 and the argument given in the last paragraph of Section 2.1 we get that the credible band (21) has honest coverage and rate adaptive size.

Corollary 1. *Assume that the eigen-vectors $\varphi_i(\cdot)$ of the linear operator $\mathbb{K}$ form the sine-cosine basis. Then there exists a constant D such that the empirical Bayes credible bands, given in* (21)*, have honest asymptotic coverage one*

$$\inf_{\theta_0 \in \Theta_s(M)} P_{\theta_0}(|\theta_0(t) - \hat{\theta}_{n, \hat{\alpha}_n - 1/2}(t)| \leq D\zeta_{\gamma/2} s_n(t, \hat{\alpha}_n - 1/2), \forall t \in [0,1])) \to 1. \tag{22}$$

Furthermore the size of the credible band is rate optimal in a minimax sense, i.e. there exists a $C_1 > 0$ such that for all $\beta \in (1/2, \beta_{\max}]$

$$\inf_{\theta_0 \in \Theta_s^\beta(M)} \mathrm{P}_{\theta_0}\big(s_n(t, \hat{\alpha}_n - 1/2) \leq C_1 n^{-\frac{\beta - 1/2}{2\beta+2p}}\big) \to 1. \tag{23}$$

3 Simulation study

We investigate our new empirical Bayes method in an example. We consider the model (1) with $\mathbb{K} = \mathbb{I}$ (the identity operator) and work with the sine-cosine basis on $[0,1]$, i.e. $\varphi_1(t) = 1$, $\varphi_{2i}(t) = \sqrt{2}\cos(2\pi it)$, $\varphi_{2i+1}(t) = \sqrt{2}\sin(2\pi it)$ for $t \in [0,1]$.

First we illustrate that for self-similar functions our method provides reasonable and trustworthy credible sets which could be used as confidence bands. We define the true function $\theta_0(t)$ with the help of its sine-cosine basis coefficients, i.e. we take $\theta_{0,i} = i^{-2}\cos(i)$:

$$\theta_0(t) = \cos(1) + \sqrt{2}\sum_{i=1}^{\infty}(2i)^{-2}\cos(2i)\cos(2\pi it) + \sqrt{2}\sum_{i=1}^{\infty}(2i+1)^{-2}\cos(2i+1)\sin(2\pi it).$$

For computational convenience we work only with the first 10^3 Fourier coefficients of the true function. We simulate data from the corresponding distribution with noise level $n = 100, 10^3, 10^4$ and 10^5. Figure 1 shows the true function in pointed black, the posterior mean in dashed red and the 95% credible bands (without blowing it up by a constant factor D) in blue. One can see that for every noise level n the credible band has good coverage and is concentrating around the truth as n increases, confirming the results of Corollary 1.

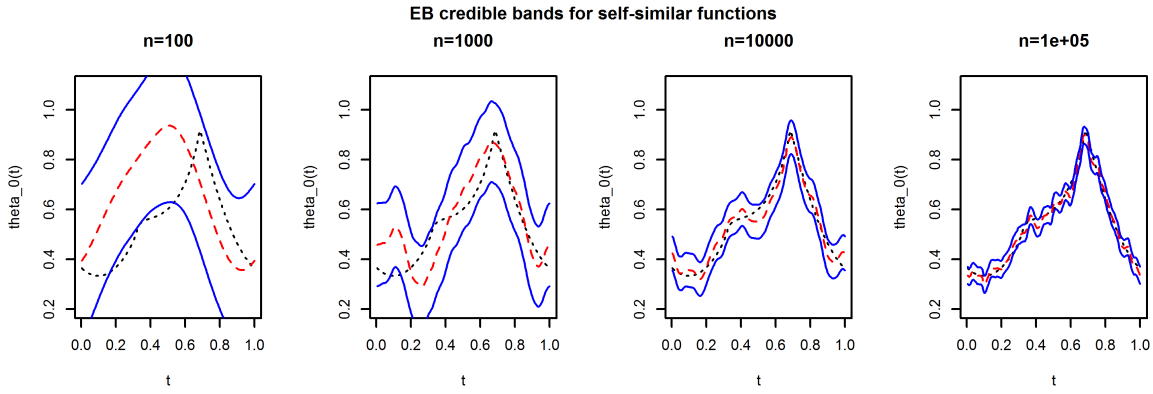

Fig. 1 Empirical Bayes credible bands for a self-similar function. The true function is drawn in pointed black, the posterior mean in dashed red and the credible bands in blue. From left to right we have $n = 100, 10^3, 10^4$ and 10^5.

To illustrate the negative result derived in Theorem 1 (for the point evaluation linear functionals) we consider a non-self-similar function $\theta_0(t)$ defined by its series decomposition coefficients with respect to the sine-cosine basis. We take the coefficients to be $\theta_{0,1} = 1/10$, $\theta_{0,4} = 1/30$, $\theta_{0,20} = -1/20$, $\theta_{0,i} = i^{-3/2}$ if $2^{4^j} < i \leq 2 * 2^{4^j}$ for $j \geq 2$, and 0 otherwise:

$$\theta_0(t) = 0.1 + \frac{\sqrt{2}}{30}\cos(4\pi t) - \frac{\sqrt{2}}{20}\cos(20\pi t) + \sum_{j=2}^{\infty}\left(\sqrt{2}\sum_{i=2^{4^{j-1}}+1}^{2^{4^j}}(2i)^{-3/2}\cos(2\pi it) + (2i+1)^{-3/2}\sin(2\pi it)\right).$$

For simplicity we consider again only the first 10^3 Fourier coefficient of the true function. Then we simulate data from the corresponding distribution with various noise levels $n = 200, 500, 10^3, 2*10^3, 5*10^3, 10^4, 10^5$ and 10^8. In Figure 2 we plotted the 95% L_∞-credible bands with blue lines, the posterior mean with dashed red line and the true function with pointed black line. One can see that for multiple noise levels we have overly confident, too narrow credible bands ($n = 500, 10^3, 2*10^3, 10^4$), while for other values of the noise levels n we have good coverage ($n = 200, 5*10^3, 5*10^4, 10^8$). These periodicity between the good and the bad coverage of the credible sets continues as n increases (but to see it we have to zoom in into the picture).

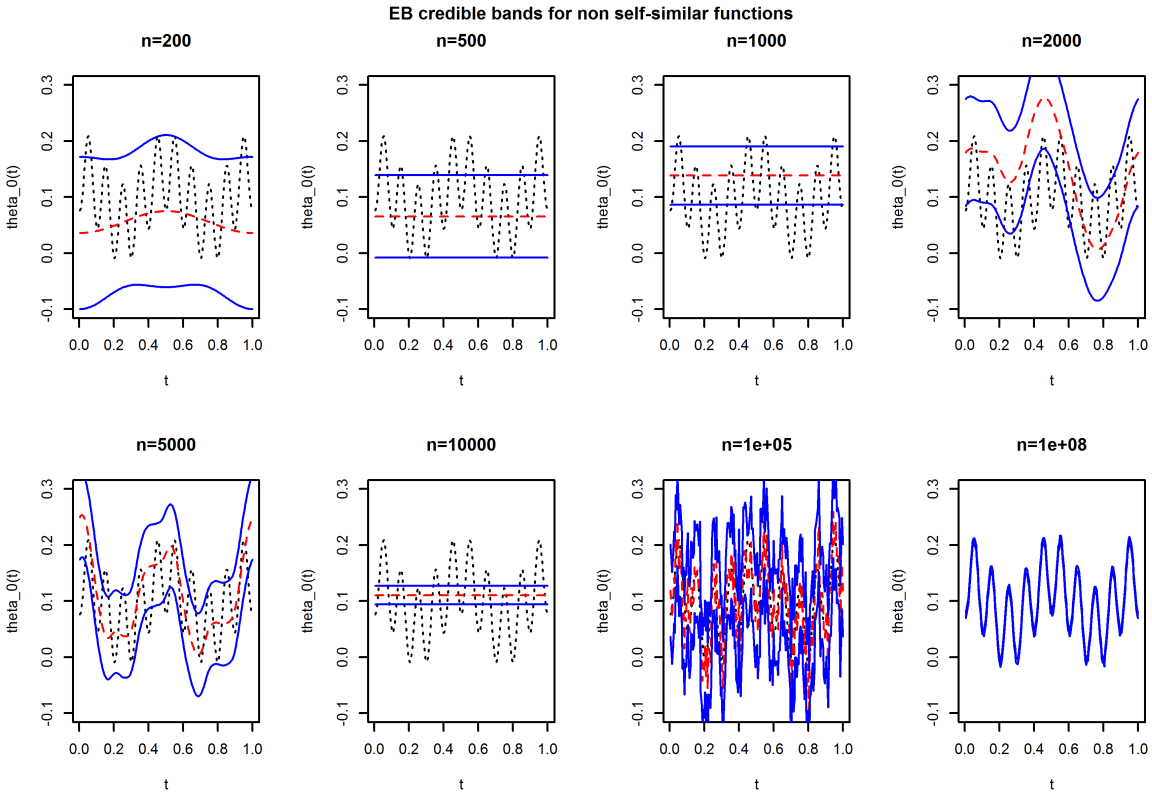

Fig. 2 Empirical Bayes credible bands for a non-self-similar function. The true function is drawn in pointed black, the posterior mean in dashed red and the credible bands in blue. From left to right and top to bottom we have $n = 200, 500, 10^3, 2*10^3, 5*10^3, 10^4, 10^5$ and 10^8.

4 Proof of Theorem 1

Following [21], [14] and [22] we introduce the notation

$$h_n(\alpha;\theta_0) = \frac{1+2\alpha+2p}{n^{1/(1+2\alpha+2p)}\log n}\sum_{i=1}^{\infty}\frac{n^2 i^{1+2\alpha}(\log i)\theta_{0,i}^2}{(i^{1+2\alpha+2p}+n)^2}, \qquad \alpha \geq 0,$$

and define

$$\underline{\alpha}_n(\theta_0) = \inf\{\alpha \in [0,\mathrm{A}] : h_n(\alpha;\theta_0) \geq 1/(16C^8)\},$$
$$\overline{\alpha}_n(\theta_0) = \sup\{\alpha \in [0,\mathrm{A}] : h_n(\alpha;\theta_0) \leq 8C^8\},$$

where the parameter A was introduced in (10). From the proof of Theorem 5.1 of [22] one can see that

$$\inf_{\theta_0 \in \ell_2} \mathrm{P}(\underline{\alpha}_n \leq \hat{\alpha}_n) \to 1. \tag{24}$$

Furthermore, let us introduce the notations

$$B_n^L(\alpha) = |\mathrm{E}_{\theta_0}\widehat{L\theta}_\alpha - L\theta_0| \quad \text{and} \quad V_n(\alpha) = |\widehat{L\theta}_\alpha - \mathrm{E}_{\theta_0}\widehat{L\theta}_\alpha|,$$

where $\widehat{L\theta}_\alpha$ denotes the posterior mean of the linear functional $L\theta$ for a fixed hyper-parameter $\alpha > 0$. Similarly to the proof of Theorem 3.1 of [22] we have following from the triangle inequality that $\theta_0 \in \hat{C}_n^L(D)$

implies $B_n^L(\hat{\alpha}_n - 1/2) \le V_n(\hat{\alpha}_n - 1/2) + D\zeta_{\gamma/2} s_n(\hat{\alpha}_n - 1/2)$. Therefore following from the convergence (24) we have

$$\mathrm{P}_{\theta_0}(L\theta_0 \in \hat{C}_n^L(D)) \le \mathrm{P}_{\theta_0}\Big(\inf_{\alpha \ge \underline{\alpha}_n - 1/2} B_n(\alpha) \le D \sup_{\alpha \ge \underline{\alpha}_n - 1/2} \big[D\zeta_{\frac{\gamma}{2}} s_n(\alpha) + V_n(\alpha) \big] \Big) + o(1). \tag{25}$$

From the proof of Theorem 3.1 of [22] follows that $\underline{\alpha}_{n_j} > \beta + 1/2$ for j large enough. Following from the proof of Theorem 5.3 of [15] we get that both $s_{n_j}(\alpha)$ and $V_{n_j}(\alpha)$ are bounded from above by a multiple of $n_j^{-(1/2+\beta+q)/(1+2\beta+2p)} \ll n_j^{-(\beta+q)/(2\beta+2p)}$ with probability tending to one.

Furthermore, for fixed hyper-parameter α the bias corresponding to the posterior mean (12) is

$$B_n^L(\alpha) = |\sum_i \frac{l_i \theta_{0,i}}{1 + n_j i^{-1-2\alpha} \kappa_i^2}|.$$

Note that following from the definition of $L_s^q(R)$ given in (5) we have that

$$\sum_{i=j}^{j+K-1} |l_i| \ge \max_{i \in \{j, j+1, \ldots, j+K-1\}} |l_i| \ge j^{-1/2-q}/(RK).$$

Then for $\alpha \ge \underline{\alpha}_{n_j} - 1/2 \ge \beta$ the squared bias $B_n^2(\alpha)$ corresponding to the sequence (15) can be bounded from below by

$$\begin{aligned} \Big(\sum_{i=n_j^{\frac{1}{1+2\beta+2p}}}^{2n_j^{\frac{1}{1+2\beta+2p}}} \frac{C_0^{-1} |l_i| K n_j^{-\frac{1/2+\beta}{1+2\beta+2p}}}{1 + n_j i^{-1-2\alpha} \kappa_i^2} \Big)^2 &\gtrsim n_j^{-\frac{1/2+\beta}{1+2\beta+2p}} \sum_{i=n_j^{1/(1+2\beta+2p)}/K}^{2n_j^{1/(1+2\beta+2p)}/K-1} \sum_{j=iK+1}^{(i+1)K} |l_i| \\ &\gtrsim n_j^{-\frac{1/2+\beta}{1+2\beta+2p}} \sum_{i=n_j^{1/(1+2\beta+2p)}/K}^{2n_j^{1/(1+2\beta+2p)}/K-1} i^{-1/2-q}, \end{aligned} \tag{26}$$

which is further bounded from below by $n_j^{-(2\beta+2q)/(1+2\alpha+2p)} \gg n_j^{-2(\beta+q)/(2\beta+2p)}$. Therefore the probability on the right hand side of the inequality (25) tends to zero.

Finally we note that following from the sub-optimal order of the bias term (26) the posterior distribution achieves sub-optimal contraction rate around the true value $L\theta_0$.

5 Proof of Theorem 2

First we note that following from the inequalities (6.9) and (6.10) of [22] we have for all $\beta \in [\beta_{\min}, \beta_{\max}]$ that

$$\beta - K_1/\log n \le \inf_{\theta_0 \in \Theta^\beta(M)} \underline{\alpha}_n(\theta_0) \le \sup_{\theta_0 \in \Theta_s^\beta(M)} \overline{\alpha}_n(\theta_0) \le \beta + K_2/\log n, \tag{27}$$

for some positive constants K_1 and K_2 depending only on $\beta_{\min}, \beta_{\max}, M, C, \rho$ and ε.

Then for convenience we introduce the notation

$$r_{n,\gamma}^2(\alpha) = D^2 \zeta_{\gamma/2}^2 s_n^2(\alpha) = D^2 \zeta_{\gamma/2}^2 \sum_{i=1}^{\infty} \frac{l_i^2 \kappa_i^{-2}}{i^{1+2\alpha} \kappa_i^{-2} + n}. \tag{28}$$

Using the notations of Section 4 the coverage of the empirical Bayes credible set, similarly to the inequality (25), can be bounded from below by

$$\inf_{\theta_0\in\Theta_s(M)} \mathrm{P}_{\theta_0}\Big(L\theta_0 \in \hat{C}_n^L(D), \forall L\in L_s^q(R)\Big) \tag{29}$$
$$\geq \inf_{\theta_0\in\Theta_s(M)} \mathrm{P}_{\theta_0}\Big(B_n^L(\hat{\alpha}_n-1/2)+V_n(\hat{\alpha}_n-1/2)\leq r_{n,\gamma}(\hat{\alpha}_n-1/2), \forall L\in L_s^q(R)\Big)$$
$$\geq \inf_{\theta_0\in\Theta_s(M)} \mathrm{P}_{\theta_0}\Big(\sup_{\substack{\alpha\in[\underline{\alpha}_n-1/2,\overline{\alpha}_n-1/2]\\ L\in L_s^q(R)}} V_n(\alpha)\leq \inf_{\substack{\alpha\in[\underline{\alpha}_n-1/2,\overline{\alpha}_n-1/2]\\ L\in L_s^q(R)}} \big[r_{n,\gamma}(\alpha)-B_n^L(\alpha)\big]\Big)-o(1).$$

Therefore it is sufficient to show that there exist constants C_1, C_2 and C_3 satisfying $C_3 > C_2 + C_1$ such that for all $\beta\in[\beta_{\min},\beta_{\max}]$ and $\theta_0\in\Theta_s^\beta(M)$

$$\sup_{\alpha\in[\underline{\alpha}_n-1/2,\overline{\alpha}_n-1/2]} B_n^L(\alpha)\leq C_1 n^{-\frac{\beta+q}{2\beta+2p}}, \tag{30}$$

$$P_{\theta_0}\Big(\sup_{\alpha\in[\underline{\alpha}_n-1/2,\overline{\alpha}_n-1/2]} V_n(\alpha)\leq C_2 n^{-\frac{\beta+q}{2\beta+2p}}\Big)\to 1, \tag{31}$$

$$\inf_{\alpha\in[\underline{\alpha}_n-1/2,\overline{\alpha}_n-1/2]} r_{n,\gamma}(\alpha)\geq C_3 n^{-\frac{\beta+q}{2\beta+2p}}. \tag{32}$$

We first deal with inequality (32). Applying assumptions (5) and (28) one can obtain that

$$r_{n,\gamma}^2(\alpha)\geq \frac{D^2\zeta_{\gamma/2}^2}{C^2}\sum_{i=j_0/K+1}^{\infty}\frac{(Ki)^{2p}}{((K+1)i)^{1+2\alpha+2p}+n}\sum_{j=iK}^{(i+1)K-1} l_j^2$$
$$\geq \frac{D^2\zeta_{\gamma/2}^2K^{2p-2q}}{R^2C^2(K+1)^{1+2\alpha+2p}}\sum_{i=j_0/K+1}^{\infty}\frac{i^{-1-2q+2p}}{i^{1+2\alpha+2p}+n},$$

which following from Lemma 1 is further bounded from below by constant times $D^2 n^{-\frac{1+2\alpha+2q}{1+2\alpha+2p}}$ for $1+2\alpha+2q>0$ and infinity else. Therefore by applying the inequality (27) we get that C_3 can be arbitrary large for a large enough choice of D^2.

Next we deal with the convergence (31). From the proof of Theorem 5.3 of [15] we get that $V_n(\alpha)=|t_n(\alpha)Z|$ with Z a standard normal random variable and

$$t_n^2(\alpha)=\sum_i^{\infty}\frac{nl_i^2\kappa_i^{-2}}{(i^{1+2\alpha}\kappa_i^{-2}+n)^2}\leq R^2C^4\sum_i^{\infty}\frac{ni^{-1-2q+2p}}{(i^{1+2\alpha+2p}+n)^2}.$$

The right hand side of the preceding display similarly to $s_n^2(\alpha)$ is bounded above by constant times $n^{-\frac{1+2\alpha+2q}{1+2\alpha+2p}}$ for $1+2\alpha+2q>0$ and infinity otherwise. Then following from the inequality (27) one can obtain for $q<p$ that

$$\sup_{\alpha\in[\underline{\alpha}_n-1/2,\overline{\alpha}_n-1/2]} t_n(\alpha)=t_n(\underline{\alpha}_n-1/2)\lesssim n^{-\frac{\underline{\alpha}_n+q}{2\underline{\alpha}_n+2p}}\lesssim n^{-\frac{\beta+q}{2\beta+2p}},$$

providing us the convergence(31).

Finally we deal with the bias term (30). Following from assumptions (3) and (5) we have

$$|B_n^L(\alpha)|\leq\sum_{i=1}^{\infty}\frac{|l_i\theta_{0,i}|i^{1+2\alpha}\kappa_i^{-2}}{i^{1+2\alpha}\kappa_i^{-2}+n}\leq C^2RK\sum_{i=1}^{\infty}\frac{i^{2\alpha+2p-\beta-q}}{i^{1+2\alpha+2p}+n}.$$

From the inequality (27) we have for $\alpha\geq\underline{\alpha}_n-1/2$ and large enough n that the inequality $\beta+q<1+2\alpha+2p$ holds, hence the preceding inequality is further bounded from above by constant times $n^{-\frac{\beta+q}{1+2\alpha+2p}}$ by applying Lemma 1 (with $m=0$). So we can conclude that for $\alpha\geq\underline{\alpha}_n-1/2\geq\beta-1/2-K_1/\log n$

$$|B_n^L(\alpha)| \lesssim n^{-(\beta+q)/(1+2\alpha+2p)} \lesssim n^{-(\beta+q)/(2\beta+2p)}.$$

To prove adaptivity we note that following again from the inequality (27) we have

$$\sup_{\alpha\in[\underline{\alpha}_n-1/2,\overline{\alpha}_n-1/2]} s_n(\alpha) \lesssim n^{-(\underline{\alpha}_n+q)(2\underline{\alpha}_n+2p)} \lesssim n^{-(\beta+q)/(2\beta+2p)}.$$

6 Appendix

Lemma 1 (Lemma 10.2 of [22]). *For any $l,m,r,s \geq 0$ with $c := lr-s-1 > 0$ and $n \geq e^{(2mr/c)\vee r}$,*

$$(3^r+1)^{-l}\big(\log n/r\big)^m n^{-c/r} \leq \sum_{i=1}^{\infty} \frac{i^s(\log i)^m}{(i^r+n)^l} \leq (3+2c^{-1})\big(\log n/r\big)^m n^{-c/r}.$$

Acknowledgements The author would like to thank Harry van Zanten, Bartek Knapik and an anonymous referee for their suggestions leading to better readability and presentation of the manuscript.

References

1. BELITSER, E. On coverage and oracle radial rate of DDM-credible sets under excessive bias restriction. *ArXiv e-prints* (July 2014).
2. BICKEL, P. J., AND KLEIJN, B. J. K. The semiparametric bernsteinvon mises theorem. *The Annals of Statistics 40*, 1 (02 2012), 206–237.
3. BULL, A. Honest adaptive confidence bands and self-similar functions. *Electron. J. Statist. 6* (2012), 1490–1516.
4. CASTILLO, I. Lower bounds for posterior rates with Gaussian process priors. *Electron. J. Stat. 2* (2008), 1281–1299.
5. CASTILLO, I. A semiparametric bernsteinvon mises theorem for gaussian process priors. *Probability Theory and Related Fields 152*, 1-2 (2012), 53–99.
6. CASTILLO, I., AND NICKL, R. On the Bernstein von Mises phenomenon for nonparametric bayes procedures. *Ann. Statist. 42*, 5 (10 2014), 1941–1969.
7. CAVALIER, L. Nonparametric statistical inverse problems. *Inverse Problems 24*, 3 (2008), 034004, 19.
8. CHERNOZHUKOV, V., CHETVERIKOV, D., AND KATO, K. Anti-concentration and honest adaptive confidence bands. *Ann. Statist.* (2014), to appear.
9. DONOHO, D. L. Statistical estimation and optimal recovery. *Ann. Statist. 22*, 1 (1994), 238–270.
10. DONOHO, D. L., LIU, R. C., AND MACGIBBON, B. Minimax risk over hyperrectangles, and implications. *Ann. Statist. 18*, 3 (1990), 1416–1437.
11. DONOHO, D. L., AND LOW, M. G. Renormalization exponents and optimal pointwise rates of convergence. *The Annals of Statistics 20*, 2 (1992), pp. 944–970.
12. FREEDMAN, D. On the Bernstein-von Mises theorem with infinite-dimensional parameters. *Ann. Statist. 27*, 4 (1999), 1119–1140.
13. GINÉ, E., AND NICKL, R. Confidence bands in density estimation. *Ann. Statist. 38*, 2 (2010), 1122–1170.
14. KNAPIK, B. T., SZABÓ, B. T., VAN DER VAART, A. W., AND VAN ZANTEN, J. H. Bayes procedures for adaptive inference in inverse problems for the white noise model. *ArXiv e-prints* (Sept. 2012).
15. KNAPIK, B. T., VAN DER VAART, A. W., AND VAN ZANTEN, J. H. Bayesian inverse problems with Gaussian priors. *Ann. Statist. 39*, 5 (2011), 2626–2657.
16. NICKL, R., AND SZABÓ, B. A sharp adaptive confidence ball for self-similar functions. *ArXiv e-prints* (June 2014).
17. PICARD, D., AND TRIBOULEY, K. Adaptive confidence interval for pointwise curve estimation. *Ann. Statist. 28*, 1 (2000), 298–335.
18. RAY, K. Bernstein-von Mises theorems for adaptive Bayesian nonparametric procedures. *ArXiv e-prints* (July 2014).
19. RIVOIRARD, V., AND ROUSSEAU, J. Bernstein–von Mises theorem for linear functionals of the density. *Ann. Statist. 40*, 3 (2012), 1489–1523.
20. SERRA, P., AND KRIVOBOKOVA, T. Adaptive empirical Bayesian smoothing splines. *ArXiv e-prints* (Nov. 2014).
21. SZABO, B. T., VAN DER VAART, A. W., AND VAN ZANTEN, J. H. Empirical bayes scaling of gaussian priors in the white noise model. *Electron. J. Statist. 7* (2013), 991–1018.
22. SZABO, B. T., VAN DER VAART, A. W., AND VAN ZANTEN, J. H. Frequentist coverage of adaptive nonparametric Bayesian credible sets. *To appear in Ann. Statist.* (2014).
23. SZABO, B. T., VAN DER VAART, A. W., AND VAN ZANTEN, J. H. Honest Bayesian confidence sets for the L2-norm. *To appear in Stat. Plann. & Inf.* (2014).
24. WEIMIN YOO, W., AND GHOSAL, S. Supremum Norm Posterior Contraction and Credible Sets for Nonparametric Multivariate Regression. *ArXiv e-prints* (Nov. 2014).